VARIANCE MINIMIZATION AND THE THEORY OF

INFLATION HEDGING

Zvi Bodie

WP 806-75   September 1975

VARIANCE MINIMIZATION AND THE THEORY OF

INFLATION HEDGING

Zvi Bodie

WP 806-75 September 1975

<u>Variance Minimization and the Theory</u>

<u>of Inflation Hedging</u>

I. Introduction

One of the casualties of the unprecedented infla-
tionary experience in this country in the past several
years has been the so-called riskless asset of portfolio
theory. The widespread practice of dismissing short-run
inflation-risk as being negligibly small and treating
short term government securities as risk-free has become
questionable, to say the least.

Although in principle it is possible to eliminate
inflation uncertainty either by creating futures markets
or by linking deferred payments to price indices,[1] no
such option is currently available in the developed
capital markets of the world.[2] In other words, at least
in these capital markets there is no perfect inflation
hedge.

In the absence of a perfect inflation hedge it is
important to determine how and to what extent one can
hedge against inflation with the financial instruments
currently available, and that is the subject of this
paper.

The term hedging against inflation as used in this paper, means reducing the specific kind of risk which stems from uncertainty about the future level of the prices of consumption goods. In other words, we use the term in exactly the same way one would use it to describe the forward purchase or sale of commodities or foreign currencies for the purpose of eliminating the risk of unanticipated changes in spot prices or exchange rates. Just as futures contracts for commodities and foreign currencies are perfect hedges against unantici- pated fluctuations in spot prices and exchange rates, so futures contracts for the specific basket of consumer goods used to define the "real" value of money (i.e., the "purchasing power" of money) would be a perfect hedge against inflation risk.

Recently there has been a revival in this country of the proposal to link deferred payments to some index of the cost-of-living like the CPI.[3] Effectively an index-linked bond is equivalent to an ordinary nominal bond plus a futures contract on the CPI. For example, imagine an index-linked Treasury Bill with a maturity of

one year and a face value of $1000 of today's purchasing power. If an investor were to purchase such a bond he would in effect be buying forward $1000 of purchasing power, but instead of contracting to pay for it in the future, which is the practice in a standard futures contract, he would be paying for it now.

The example of an index-linked Treasury Bill can also be used to clarify a very important point about hedging against inflation which is often overlooked. If the Treasury were in fact to issue the 1-year Bill described above it might under current market conditions be able to sell it at a price in excess of $1000. In other words the riskless real rate of interest on such a bond could very well be negative.

Now many people would claim that if it has a negative real rate of return then the security is not an inflation hedge. In their view an inflation hedge is an asset which "keeps up" with inflation, i.e., has a nominal rate of return at least as great as the rate of inflation.

It does not take extensive research to discover that no such security currently exists in the financial markets of this country and that there is no way for an investor to create a portfolio having that feature

(i.e., a guaranteed non-negative real rate of return).
Furthermore, it should be added that the proponents of
issuance by the Federal Government of index-linked
bonds have not stipulated that the real rate of interest
on these bonds be non-negative, although it would probably
be fair to say that they have assumed that the rate at
least on longer maturities would be positive.[4]

The point is that there is nothing in the definition
of an inflation hedge as used in this study which would
imply that it has a non-negative real rate of return or
that an investor would want to hold it. Although it is
natural to assume that risk-averse consumer-investors
would want to hedge some of their inflation risks, how
much to hedge would depend among other things on the
cost of hedging.

In order to deal with the question of hedging against
inflation the issues must be more sharply defined. Since
there is one type of security whose real return is
certain but for inflation risk, namely single-period,
riskless-in-terms-of default money-fixed bonds, it seems
natural to identify inflation risk with the stochastic
component of the real return on such a bond. Therefore
in this paper the conceptual criterion for deciding the

extent to which a particular security or class of
securities is an inflation hedge is the extent to
which it can be used to reduce the uncertainty of
the real return on a nominal bond.

We adopt the convention of identifying the
riskiness of a probability distribution with its
variance. Accordingly, we measure the effectiveness
of a security as an inflation hedge as the proportional
reduction in the variance of the real return on a one-
period, risk-free-in-terms-of-default, nominal bond
attainable by combining the hedge security and the bond
in their variance minimizing proportions.

It is worthwhile to indicate at this point the
relationship between this view of hedging against infla-
tion and the investor's ultimate objective of optimal
portfolio selection. This can best be done in the frame-
work of the Markowitz-Tobin mean-variance model of
portfolio choice.[5]

In that model the process of portfolio selection
is divided into two separate stages: (1) identification
of the efficient portfolio frontier and (2) choosing
the optimal portfolio on that frontier. This paper
focuses on one particular point on the efficient
frontier -- the minimum variance portfolio. From this
perspective hedging against inflation is essentially the
process of taking a risk-free-in-terms-of-default nominal
bond as the starting point and using other securities to
eliminate as much of the variance of its real return as
possible.

We define the difference between the mean real
return on a nominal bond and the mean real return on
the minimum variance portfolio as the "cost" of hedging
against inflation. Since the unhedged nominal bond does
not in general lie on the efficient portfolio frontier,
the cost of hedging may be either positive or negative.

The approach to hedging against inflation adopted
in this paper can perhaps be further clarified with the
aid of figure 1, in which mean real return is measured
along the vertical axis and standard deviation of real
return along the horizontal. Point A is the hypothetical
location of an unhedged nominal bond. The minimum
variance portfolio lies somewhere to the left of A

Figure 1

(e.g., points B or C); the farther to the left it is, the more effective the inflation hedge. The cost of hedging is represented in the diagram by the vertical distance between point A and the point representing the minimum variance portfolio. Thus if point B were the minimum variance portfolio the cost of hedging would be positive; if point C, the cost of hedging would be negative.

The minimum-variance portfolio is in principle composed of all available securities. Assume that there are $n + 1$ such securities, with the $n + 1$st being nominal bonds. In part II of the paper we view the minimum variance portfolio as being composed of just two securities: the nominal bond and the portfolio of the other n securities, which we refer to as the optimal inflation-hedge security, or simply the hedge security. Part II tries to relate the effectiveness and the cost of hedging to the parameters of the joint probability distribution of the real returns on the two component securities.

Part III deals with the composition of the optimal
inflation-hedge portfolio, while part IV tries to
indicate how uncertainty about the parameters of the
joint probability distribution of security returns
affects one's ability to hedge against inflation.
In part V we try to relate the theory of inflation
hedging developed in the previous parts of the paper
to modern portfolio theory, and specifically, to the
capital asset pricing model. Finally, in part VI we
summarize the results of the paper and draw its
major conclusions.

II. The Determinants of Hedging Effectiveness.

Let $P(t)$ be the price level at time t, i.e., the nominal value at time t of some standard basket of commodities. Let $r_i(t)$ be the continuously compounded one-period nominal rate of return on security i during period t. The continuously compounded real rate of return on security i is then given by:

$$\tilde{R}_i(t) \equiv \tilde{r}_i(t) - \tilde{\pi}(t)$$

where $\pi(t)$ is the rate of inflation during period t defined by:

$$e^{\tilde{\pi}(t)} = \frac{\tilde{P}(t+1)}{P(t)}$$

and tildes (~) are used to denote random variables.

Suppose the security under consideration is a bond that matures at time t + 1 and is free of default risk, so that its nominal rate of return, $r_n(t)$, is known with certainty at time t. Its real return, however, is uncertain at time t, since it depends on $P(t+1)$.

Let us decompose $\tilde{\pi}(t)$ into its mean, $\bar{\pi}(t)$, and deviation from its mean, $p(t)$. $\tilde{p}(t)$ represents the "unanticipated" rate of inflation. Suppressing the time subscripts, the real rate of return on a risk-free nominal bond is simply:

(1) $\tilde{R}_n \quad r_n - \pi - \tilde{p}$

or $\tilde{R}_n \quad \bar{R}_n - \tilde{p}$

Now consider the second security, the portfolio of risky securities used to hedge the nominal bond against inflation and hereafter referred to as the hedge security. Let us also decompose its real rate of return into its mean and deviation from the mean:

$$(2) \qquad \tilde{R}_h = \bar{R}_h + \tilde{\epsilon}$$

The deviation from the mean, $\tilde{\epsilon}$, can be further decomposed into two orthogonal components as follows:[6]

$$\tilde{\epsilon} = \alpha_R \tilde{p} + \tilde{\mu}$$

where $\qquad \alpha_R = \dfrac{\text{cov}(\epsilon, p)}{\text{var}(p)} \qquad$. (2) then becomes:

$$(3) \qquad \tilde{R}_h \equiv \bar{R}_h + \alpha_R \tilde{p} + \tilde{\mu}$$

$\alpha_R \tilde{p}$ can be called the "inflation-risk" component of the real return on the hedge security while $\tilde{\mu}$ is the "non-inflation risk" component.

It should be stressed that since these relationships are definitionally true, they assume nothing and they imply nothing about the causal link or the time series relationship between the return on the security and the rate of inflation. Neither do they assume anything about the specific probability distributions of $\tilde{p}$ and $\tilde{\mu}$.

Now if we replace the real rate of return on the hedge security with the nominal rate of return then equation (3) becomes:

$$(3') \quad r_h \equiv \bar{r}_h + \alpha \tilde{p} + \tilde{\mu}$$

where $\alpha \equiv \alpha_R + 1$

Our objective is to combine nominal bonds and the hedge security so as to create the portfolio with minimum variance of real return. Letting w be the proportion of the hedge security in the minimum variance portfolio, the latter's real rate of return, $\tilde{R}_{min}$ is to a very close approximation:[7]

$$\tilde{R}_{min} = \tilde{R}_n + w(\tilde{R}_h - \tilde{R}_n)$$

or by substituting from (1) and (3):

(4) $\quad \tilde{R}_{min} = \bar{R}_n + w(\bar{R}_h - \bar{R}_n) + (w\,\alpha-1)\tilde{p} + w\mu$

Let σ_1^2 be the variance of p, σ_2^2 the variance of μ, and σ_{min}^2 the variance of $\bar{R}_{min}$.

From (4) we get:

(5) $\quad \sigma_{min}^2 = (w\,\alpha-1)^2 \sigma_1^2 + w^2 \sigma_2^2$

The value of w which minimizes σ_{min}^2 is:

(6) $\qquad w = \cfrac{\alpha}{\alpha^2 + \cfrac{\sigma_2^2}{\sigma_1^2}}$

Substituting from (6) into (5), the value of the minimized variance is given by:

(7) $\qquad \sigma_{min}^2 = \cfrac{1}{\cfrac{\alpha^2}{\sigma_2^2} + \cfrac{1}{\sigma_1^2}}$

Our measure of the effectiveness of the hedge security
as an inflation hedge is the proportional reduction in the
variance of the real return on the nominal bond attainable
by combining the two securities in the variance minimizing
proportions. The formula for this measure is:

$$
(8) \qquad S = \frac{1}{1 + \dfrac{\sigma_2^2}{\sigma_1^2 \, \alpha^2}}
$$

It is a simple matter to show that S is none other than
ρ^2, the coefficient of determination between r_h and $\tilde{p}$.[8]
Thus one can say that the more highly correlated (either
negatively or positively) the nominal rate of return on a
security is with unanticipated inflation, the more effective
it is as an inflation hedge. Note that if the investor
cannot sell the hedge security short then it is an inflation
hedge if and only if its nominal return is positively corre-
lated with unanticipated inflation.

Finally as a measure of the "cost" of hedging we take
the difference between the expected values of the real return
on the unhedged nominal bond and the real return on the mini-
mum variance portfolio, which from (4) is:

(9) $C \equiv E(\tilde{R}_n - \tilde{R}_p) \equiv w(\bar{R}_n - \bar{R}_h)$

Equations (6) and (8) reveal that both w and ρ^2 are
functions of σ_2/σ_1 and α . σ_2/σ_1 is the ratio of the
measure of the non-inflation risk in the real return on
the hedge security to the measure of inflation risk, while
α is the regression coefficient of the nominal rate of return
on the hedge security on the unanticipated rate of inflation.

Now let us explore the implications of equations (6)
and (8) with the aid of figure 2, in which σ_2/σ_1 is measured
along the horizontal axis and α along the vertical. First
let us direct our attention to (8), the formula for ρ^2 or S.

Figure 2

Each "iso-effectiveness" curve in figure 2, i.e. each
locus of parameter values yielding the same value of ρ^2, is
a set of 2 straight lines which meet at the origin and are
symmetric with respect to the horizontal axis. There are two
important exceptions at the extremes. The iso-effectiveness
curve corresponding to an ρ^2 value of zero is the horizontal
axis itself (excluding the origin), and the iso-effectiveness
curve corresponding to an ρ^2 value of one is the vertical
axis (excluding the origin). The closer the iso-effectiveness
lines are to the vertical axis (i.e., the greater the angle
they make with the horizontal axis), the higher the value of ρ^2.

Along the horizontal axis the nominal rate of return on
the hedge security is uncorrelated with unanticipated inflation
so no reduction in variance is possible. The reason for this
can most easily be seen by comparing the expressions for $\tilde{R}_h$ and
$\tilde{R}_n$ in this case:

$$\tilde{R}_n \equiv \bar{R}_n - \tilde{p}$$
$$\tilde{R}_h = \bar{R}_h - \tilde{p} + \tilde{\mu}$$

Because the regression coefficients on $\tilde{p}$ are the same, combining
any amount of the hedge security, positive or negative, with the
nominal bond would just add non-inflation risk to the portfolio
without eliminating any of the inflation risk. Equation (5)

confirms this since when $\alpha = 0$ it reduces to:

$$\sigma^2_{min} = \sigma^2_1 + w^2\sigma^2_2$$

Similarly equation (6) indicates that when $\alpha = 0$, w is zero.

Along the vertical axis, on the other hand, ρ^2 is 1 because $\sigma_2 = 0$, i.e., the only source of uncertainty in the return on the hedge security is unanticipated inflation. In that case the rates of return on the nominal bond and the hedge security are perfectly correlated, and they can therefore be combined to create a real risk-free composite security.

For a given value of α (different from zero), ρ^2 falls monotonically as σ_2/σ_1 rises indicating that the greater the degree of non-inflation risk in the return on the hedge security relative to the degree of uncertainty about inflation, the less effective the security is as an inflation hedge. For a given value of σ_2/σ_1, on the other hand, ρ^2 rises monotonically as $|\alpha|$ (i.e., the distance from the horizontal axis) increases.

Consequently, along all but the two extreme iso-effectiveness curves there is a trade-off between $|\alpha|$ and σ_2/σ_1 in the sense that the loss of hedging effectiveness resulting from a rise in σ_2/σ_1 can be compensated by a rise in $|\alpha|$, the "marginal rate of substitution" between $|\alpha|$ and σ_2/σ_1 for a given value of ρ^2 being a constant.

Figure 2 can also be used to illustrate the implications of equation (6), the formula for w. As already stated above, if $\alpha = 0$ no reduction in variance can be achieved through diversification, and therefore w equals zero all along the horizontal axis (except at the origin where it is undefined). w has the same sign as α indicating that the hedger must take a long position if the nominal rate of return is positively correlated with unanticipated inflation and a short position if negatively correlated.

An iso-w curve in figure 2 is defined as the locus of all combinations of parameter values which yield minimum-variance portfolios containing the same proportion of the hedge security.

The iso-w curves for positive values of w form a set of ever-widening semi-circles which fan out from the origin in an upward direction. Although they all converge at the origin, the origin itself is not on any of them. The set of iso-w curves for negative values of w is symmetric to the positive set with respect to the horizontal axis, i.e. for every positive-valued iso-w curve fanning out upward from the origin there is an identical negative-valued iso-w curve with the same absolute value fanning downward from the origin.

The only iso-w curve shown in figure 2 is the one corresponding to a value of 1. For any combination of parameter values lying on this semi-circle, the minimum-variance portfolio will consist solely of the hedge security; it will contain no nominal bonds. The shaded area within the unit iso-w curve contains all the parameter combinations at which nominal bonds must be sold short in order to create the minimum variance portfolio.

The two straight lines emanating from the origin and
making 45 degree angles with the horizontal axis have special
significance. Firstly, at all points along these lines
$\rho^2 = .5$. But more importantly, if we hold σ_2/σ_1 constant and
examine the way w varies as a function of α, we find that w
reaches its maximum where $\alpha = \sigma_2/\sigma_1$ and its minimum where
$\alpha = -\sigma_2/\sigma_1$. In other words if we drew a vertical line in
figure 2 through the point on the horizontal axis corres-
ponding to some value of the standard deviation ratio, w
would be at its minimum where the line intersects the lower
45 degree line and at its maximum where the line intersects
the upper 45 degree line. Along the upper 45 degree line
$w = \dfrac{\sigma_1}{2\sigma_2}$ and along the lower one $w = -\dfrac{\sigma_1}{2\sigma_2}$.

Perhaps further light can be shed on the relationships between ρ^2 and α and between w and α by examining figure 3 in which α is measured along the horizontal axis and w along the vertical. Each curve in figure 3a shows the functional relationship between w and α for a given value of σ_2/σ_1. The effect of a decrease in σ_2/σ_1 is illustrated by the shift from curve 1, which corresponds to the higher value of σ_2/σ_1, to curve 2. Curve 2 is much steeper than curve 1 in the vicinity of the origin (i.e., for small values of $|\alpha|$) indicating that the smaller the value of σ_2/σ_1 the more sensitive is w to $|\alpha|$. Thus if σ_2/σ_1 and $|\alpha|$ are both small, the hedger must be either very long or very short the hedge security depending on the sign of α.

In figure 3b we see the graphs of the two ρ^2 functions corresponding to curves 1 and 2 in figure 3a. The smaller the value of σ_2/σ_1 the "closer" the two branches of the ρ^2-curve to the vertical axis. As in figure 3a, curve 2 is much steeper than curve 1 in the vicinity of the origin thus indicating that the smaller the value of σ_2/σ_1 the more sensitive is ρ^2 to $|\alpha|$.

Figure 3

a.

b.

In the limit, when $\sigma_2 = 0$, perfect hedging is possible so that $\rho^2 = 1$ for all values of α except zero. Curve 3 in figure 3a shows how w varies with α in this limiting case. Note that instead of passing through the origin as do curves 1 and 2, curve 3 splits into two parts, both of which are asymptotic to the vertical axis. This is because in the special case where $\sigma_2 = 0$, α cannot be equal to zero or else the hedge security would be indistinguishable from nominal bonds in the stochastic component of its real rate of return and would either dominate (if $\bar{R}_h > \bar{R}_n$) or be dominated by nominal bonds (if $\bar{R}_n > \bar{R}_h$).

Up to this point the discussion has dealt exclusively with variance reduction. Let us now examine the implications of equation (9), the formula for the cost of hedging.

Equation (9) indicates that C will be positive in either of two cases. The first is if both the mean real return on the nominal bond exceeds that of the hedge security and the hedger must take a long position in the hedge security (i.e., $\alpha > 0$). The second case is if $R_n < R_h$ and α is negative. In either case C is positive, meaning that the mean real return on the minimum variance portfolio would be less than the mean real return on the unhedged nominal bond.

C will be negative, i.e., the minimum variance portfolio will have a greater mean real return than an unhedged nominal bond, in either of the remaining two possible cases, namely:

(1) $\bar{R}_n > \bar{R}_h$ and w is negative.

(2) $\bar{R}_n < \bar{R}_h$ and w is positive.

The following table summarizes these relationships.

Cost of hedging		
	$\bar{R}_n > \bar{R}_h$	$\bar{R}_n < \bar{R}_h$
$\alpha > 0$	+	−
$\alpha < 0$	−	+

III. The Optimal Inflation-Hedge Portfolio.

Up to this point we have treated the creation of the minimum variance portfolio as a problem involving just two securities. But in fact there are $n + 1$ risky securities to choose from, the $n + 1^{st}$ being nominal bonds. The purpose of this section is to examine the composition of the portfolio of the first n securities, which is combined with nominal bonds to create the minimum variance portfolio.

Let R_i be the continuously compounded real rate of return on security i. As in section II decompose it into its mean and deviation from the mean:

$$R_i \equiv \bar{R}_i + \tilde{\epsilon}_i$$

and decompose $\tilde{\epsilon}_i$ into an inflation-risk component and a non-inflation risk component:

$$\tilde{\epsilon}_i = \alpha_i^R \tilde{p} + \tilde{\mu}_i$$

Since security $n + 1$ is nominal bonds we know that:

$$\tilde{\epsilon}_{n+1} = -\tilde{p}$$

so that $\alpha_{n+1}^R = -1$ and $\mu_{n+1} \equiv 0$.

Let σ_{ij}^R be the covariance between $\tilde{R}_i$ and $\tilde{R}_j$, σ_1^2 the variance of $\tilde{p}$, and σ_{2ij} the covariance between $\tilde{\mu}_i$ and $\tilde{\mu}_j$. Then:

$$\sigma_{ij}^R = \alpha_i^R \alpha_j^R \sigma_1^2 + \sigma_{2ij} \quad .$$

The variance of the real return on any portfolio composed of these securities is given by:

$$\sigma^2 \equiv \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} w_i w_j \sigma_{ij}$$

or

$$\sigma^2 = \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} w_i w_j (\alpha_i^R \alpha_j^R \sigma_1^2 + \sigma_{2ij})$$

and

$$\sum_{i=1}^{n+1} w_i = 1$$

Let w_i^* be the portfolio weight of security i in the minimum-variance portfolio. The weight of security i in the optimal inflation-hedge portfolio composed of the first n securities is therefore given by:

$$w_i' = \frac{w_i^*}{\sum_{j=1}^{n} w_j^*}$$

.

Note that $\sum_{i=1}^{n} w_i' = 1$.

The real rate of return on the optimal inflation-hedge portfolio is therefore:

$$\bar{R}_h = \sum_{i=1}^{n} w_i' \tilde{R}_i$$

As in equation (3) R_h can be decomposed into its mean, $\bar{R}_h$, an inflation-risk component, $\alpha^R \tilde{p}$ and a non-inflation risk component $\tilde{\mu}$.

The values of α^R and σ_2^2 for this portfolio are given by:

$$\alpha^R = \sum_{i=1}^{n} w_i' \alpha_i^R$$

$$\sigma_2^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i' w_j' \sigma_{2ij}$$

It is shown in the appendix that the composition of the optimal inflation-hedge portfolio (i.e., the set of weights w_i' i=1, n) is independent of σ_1^2, the variance of unanticipated inflation.

This is a kind of separability or mutual fund theorem for hedging against inflation because it implies that a hedger can create the minimum variance portfolio in two separate stages.

In the first stage he selects a portfolio out of all securities other than the risk-free nominal bond so as to maximize the coefficient of determination between the nominal return on the resultant portfolio and unanticipated inflation. The relative pro-portions of the various securities in this optimal inflation-hedge portfolio depend only on α_i, $i = 1..., n$, and the variance-covariance matrix of the non-inflation stochastic components of the securities' rates of return, but not on σ_1^2.

In the second stage, a judgment is made about the variance of unanticipated inflation and the optimal inflation-hedge portfolio is combined with nominal bonds to create the minimum variance portfolio as described above in section II.

Thus even though investors may have diverse judgments about the variance of unanticipated inflation, if they agree about the values of all the other parameters then they will agree about the composition of the optimal inflation-hedge portfolio to be combined with nominal bonds to form the minimum-variance portfolio.

IV. The Effect of Parameter Uncertainty on Hedging Effectiveness

Up to this point we have assumed that α, σ_1^2 and σ_2^2 are known. Generally, however, they will themselves be subject to uncertainty. In that case the portfolio proportion actually chosen for the hedge security will not necessarily be equal to the true variance-minimizing proportion. Let w be the proportion actually chosen and w the "correct" proportion. Define:

$$\varepsilon = \hat{w} - w$$

From (5) we have:

$$(13) \quad \hat{\sigma}^2 = \sigma_1^2 - 2\hat{w}\alpha\sigma_1^2 + \hat{w}^2(\alpha^2\sigma_1^2 + \sigma_2^2)$$

where $\hat{\sigma}^2$ is the variance of the real return on the portfolio actually selected. By substituting $w + \varepsilon$ for $\hat{w}$ in (13) we get:

$$\hat{\sigma}^2 = \sigma_1^2 - 2w\alpha\sigma_1^2 + w^2(\alpha^2\sigma_1^2 + \sigma_2^2) - 2\varepsilon\alpha\sigma_1^2 + (\varepsilon^2 + 2w\varepsilon)(\alpha^2\sigma_1^2 + \sigma_2^2)$$

which by substituting for w from (6) reduces to:

$$(14) \quad \hat{\sigma}^2 = \sigma_1^2 - \left(\frac{1}{1 + \dfrac{\sigma_2^2}{\alpha^2\sigma_1^2}}\right)\sigma_1^2 + \varepsilon^2(\alpha^2\sigma_1^2 + \sigma_2^2)$$

The first term in (14) is the variance of the unhedged nominal bond; the second term is the reduction in that variance attainable by hedging the bond against inflation assuming that

one knows the correct variance-minimizing proportion for the
hedge security; and the third term is the increase in variance
due to using a proportion for the hedge security which differs
by ϵ from the correct variance-minimizing proportion. If the
third term is greater than the absolute value of the second,
the investor is really increasing the variance of his real
return by trying to hedge against inflation.

To gain further insight into the effect of parameter un-
certainty let us assume that the source of the error in the
portfolio proportion is an error in estimating α. Thus let:

$$\hat{\alpha} \equiv \alpha + \eta$$

where $\hat{\alpha}$ is our estimate of α and η is the deviation of $\hat{\alpha}$ from
the "true" value.

In that case the deviation of our chosen portfolio propor-
tion for the hedge security from the correct variance-minimizing
proportion will be:

$$(15) \qquad \epsilon = \sigma_1^2 \left\{ \frac{\eta \sigma_h^2 - \alpha(\eta^2 + 2\alpha\eta)\sigma^2}{\sigma_h^2[\sigma_h^2 + (\eta^2 + 2\alpha\eta)\sigma_1^2]} \right\}$$

where we let σ_h^2 stand for the variance of the nominal return
on the hedge security and make use of the identity:

$$\sigma_h^2 = \alpha^2\sigma_1^2 + \sigma_2^2$$

Now suppose that $\alpha = 0$, i.e., the hedge security is un-correlated with unanticipated inflation, but we mistakenly assume that the regression coefficient is not zero. In this case (15) reduces to:

(16)
$$c = \frac{\eta \sigma_1^2}{\sigma_2^2 + \eta^2 \sigma_1^2}$$

which is also the formula for $\hat{w}$.

Substituting into (14) we get:

(17)
$$\hat{\sigma}^2 = \left[1 + \frac{\eta^2 \sigma_1^2 \sigma_2^2}{(\eta^2 \sigma_1^2 + \sigma_2^2)^2} \right] \sigma_1^2$$

Equation (17) implies that the "damage" done by not leaving well enough alone (i.e., by combining some of the mistaken hedge security with the risk-free nominal bond) is a mono-tonically increasing function of η^2, σ_1^2 and σ_2^2. In practice if α is estimated using standard regression techniques then η^2, the square of the error in the estimate of α, will be larger the larger the value of σ_2^2 (i.e., the residual variance). A reasonable conclusion to draw from this discussion, therefore, is that a security whose nominal return has a lot of non-inflation risk in it should probably not be used as an inflation hedge even if the point estimate of its α value is quite dif-ferent from zero.

V.   Inflation-risk, Market-risk and Equilibrium Returns.

In modern capital market theory a basic distinction is drawn between systematic and unsystematic risk. In the Sharpe-Lintner-Mossin version of the capital asset pricing model (CAPM) [9] the two most important theorems are:

(1)   The equilibrium expected excess return on a security
       is proportional to its systematic risk as measured by
       the regression coefficient of its return on the return
       of the "market portfolio."

(2) Efficient portfolios have no unsystematic risk.

In the absence of a riskless asset, the first of these propositions is still correct if one measures excess returns with reference to the minimum-variance-zero-beta portfolio instead of the riskless asset. [10]

Since a large proportion of the wealth of the private sector is in the form of money-fixed assets there is undoubtedly a negative correlation between unanticipated inflation and the return on the market portfolio, [11] so that the theory would predict a positive risk premium (above the return on the minimum variance zero beta portfolio) on any security negatively correlated with unanticipated inflation (i.e., positively correlated with the real return on money-fixed assets).

Unlike the first theorem of the CAPM, the second does not survive the transition to a capital market with no riskless asset. In such a market efficient portfolios will in general have some and perhaps a great deal of unsystematic risk.

Let $\tilde{R}_m(t)$ represent the real return on the market portfolio in period t and decompose it into its mean and deviation from the mean:

$$\tilde{R}_m(t) \equiv \overline{R}_m(t) + \tilde{m}(t)$$

The real return on security i can then be decomposed into its mean plus the sum of two orthogonal stochastic components:

$$\tilde{R}_i(t) \equiv \overline{R}_i(t) + \beta_i(t)\tilde{m}(t) + \tilde{\eta}_i(t)$$

where $\beta_i(t) \equiv \dfrac{cov(\tilde{R}_i(t),\tilde{m}(t))}{var(\tilde{m}(t))}$. The first, $\beta_i(t)\tilde{m}(t)$, is the systematic risk of security i, and $\tilde{\eta}_i(t)$ is its unsystematic risk.

By the first theorem of the CAPM:

$$E(\tilde{R}_i(t)) = E(\tilde{R}_0(t)) + \beta_i(t)\ E(\tilde{R}_m(t) - \tilde{R}_0(t))$$

where $E(\ )$ is the expectation operator and $\tilde{R}_0$ is the real rate of return on the minimum-variance zero-beta portfolio. If, as seems likely, the market portfolio is negatively correlated with p(t), the unanticipated change in the rate of inflation, then the value of $\beta_i(t)$ for an unhedged nominal bond is positive implying a positive risk premium for $\tilde{R}_n$.

Since the real returns on all efficient portfolios are positively correlated, [12] the minimum variance portfolio also has a positive risk premium. Which security has a higher return, and therefore whether the "cost" of hedging is positive or negative, cannot be determined on the basis of purely theoretical considerations.

Note, however, that if the market portfolio is uncorrelated or positively correlated with inflation then a nominal bond will have a non-positive beta. In that case the expected return on the minimum variance portfolio will exceed the expected return on an unhedged nominal bond and therefore the cost of hedging will be negative.

Although these theoretical considerations do not tell us whether the cost of hedging is positive or negative they do impose an upper bound on the mean real return on a zero-beta portfolio. Specifically, if the theory is correct then the mean real return on a zero-beta portfolio must be less than the expected return on the minimum variance portfolio.

VI.   Summary and Conclusions

The point of departure for this paper was the view that inflation hedging is essentially the process of creating the minimum variance portfolio by taking a risk-free-in-terms-of default money-fixed bond and using other securities to eliminate as much of its variance as possible.

In part II it was shown that the extent to which this variance, which was identified with inflation-risk, can be reduced depends on the coefficient of determination between the rate of return on the hedge security and unanticipated inflation. The higher the value of this parameter, the larger the reduction in variance. Thus assuming short sales are possible one can say that a security is an inflation hedge to the extent that its nominal rate of return is correlated - either positively or negatively - with unanticipated inflation.

The coefficient of determination was then broken down and analyzed in terms of two other parameters. The first of these was the ratio of the variance of the non-inflation stoch-astic component of the rate of return on the hedge security to the variance of unanticipated inflation. The larger this variance ratio, the smaller the coefficient of determination. The second parameter was the coefficient of unanticipated inflation in the regression equation for the nominal rate of

return on the hedge security. The greater the absolute value of this regression   coefficient, the greater the coefficient of determination.

In part III we proved a separability or mutual fund theorem for hedging against inflation. It was shown that even though investors may have diverse judgements about the variance of unanticipated inflation, if they agree about the values of all other parameters then they will agree about the composition of the optimal inflation-hedge portfolio to be combined with nominal bonds to form the minimum-variance portfolio.

In part IV we showed how uncertainty about the parameters of the joint probability distribution of security returns might lead to the paradoxical situation where the would-be hedger is actually increasing rather than decreasing the variance of his real return.

Finally, in part V we showed that although portfolio theory does have some implications for the cost of hedging against inflation, it is impossible to determine on the basis of theoretical considerations alone whether that cost is positive or negative.

Appendix

Proof that w_i' $i = 1,\ldots, n$ do not depend on σ_1^2.

By assumption we hold α_1 and σ_{2ij} fixed for the first n

securities and consider the effect of changing the value of σ_1^2.

Let w_i' represent the weight of security i in the optimal

inflation-hedge portfolio corresponding to a value for σ_1 of

$\hat{\sigma}_1$ and w_i'' be the weight corresponding to a value of $\hat{\hat{\sigma}}_1$,

where $\hat{\sigma}_1 \neq \hat{\hat{\sigma}}_1$. From equation (7) we have:

$$\hat{\sigma}_{min}^2 = \cfrac{1}{\cfrac{\hat{\alpha}^2}{\hat{\sigma}_2^2} + \cfrac{1}{\hat{\sigma}_1^2}} \quad \text{and} \quad \hat{\hat{\sigma}}_{min}^2 = \cfrac{1}{\cfrac{\hat{\hat{\alpha}}^2}{\hat{\hat{\sigma}}_2^2} + \cfrac{1}{\hat{\hat{\sigma}}_1^2}}$$

where:

$$\hat{\alpha} = \sum_{i=1}^{n} w_i'\alpha_1 \qquad \hat{\hat{\alpha}} = \sum_{i=1}^{n} w_i''\alpha_1$$

$$\hat{\sigma}_2^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i'w_j'\sigma_{2ij} \qquad \hat{\hat{\sigma}}_2^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i''w_j''\sigma_{2ij}$$

From the fact that $\hat{\sigma}_{min}^2$ is a minimum for $\sigma_1 = \hat{\sigma}_1$ we

have:

$$\hat{\sigma}_{min}^2 = \cfrac{1}{\cfrac{\hat{\alpha}^2}{\hat{\sigma}_2^2} + \cfrac{1}{\hat{\sigma}_1^2}} \leq \cfrac{1}{\cfrac{\hat{\hat{\alpha}}^2}{\hat{\sigma}_2^2} + \cfrac{1}{\hat{\sigma}_1^2}}$$

which implies that:

(1)
$$\frac{\hat{\alpha}^2}{\hat{\sigma}_2^2} \geq \frac{\hat{\hat{\alpha}}^2}{\hat{\hat{\sigma}}_2^2}$$

Similarly,

$$\hat{\hat{\sigma}}^2_{min} = \frac{1}{\dfrac{\hat{\hat{\alpha}}^2}{\hat{\sigma}_2^2} + \dfrac{1}{\hat{\sigma}_1^2}} \leq \frac{1}{\dfrac{\hat{\alpha}^2}{\hat{\sigma}_2^2} + \dfrac{1}{\hat{\sigma}_1^2}}$$

which implies that:

(2)
$$\frac{\hat{\alpha}^2}{\hat{\sigma}_2^2} \geq \frac{\hat{\hat{\alpha}}^2}{\hat{\sigma}_2^2}$$

From (1) and (2) we have:

(3)
$$\frac{\hat{\alpha}^2}{\hat{\sigma}_2^2} = \frac{\hat{\hat{\alpha}}^2}{\hat{\hat{\sigma}}_2^2}$$

Since (3) must hold for all possible values of α_i and $\sigma_{2_{ij}}$ it must be that:

$$w'_i = w''_i \qquad i = 1,\dots,n$$

Q.E.D.

Footnotes

[1] This paper ignores the ambiguities and difficulties involved in defining and measuring the general price level. It assumes that the Consumer Price Index or some such index is an appropriate measure.

[2] Although there are futures markets for many basic commodities, these do not generally correspond to the components of a basket of final consumption goods. The idea of establishing futures markets for the all-item CPI and its components has been suggested by Lovell and Vogel (1971).

[3] See Bhatia (1974) for a review of this literature.

[4] See the articles by Friedman (1974), Modigliani (1974), and Tobin (1963) for example.

[5] Originally formulated by Markowitz (1952), this model was thought to be consistent with expected utility maximization only under very restrictive assumptions about either the stochastic specification or the utility function. Recent work by Merton, (1969, 1971) however, has shown that the mean-variance model has validity for a broad class of stochastic specifications and utility functions when the trading interval is sufficiently small.

[6] Proof that $\bar{p}$ and $\tilde{p}$ are orthogonal, i.e., that their co-variance is identically zero:

$$cov(\bar{p}, \tilde{p}) = cov(\bar{p}, r - \alpha_R \bar{p})$$

$$= cov(\bar{p}, \tilde{r}) - \alpha_R var(\bar{p})$$

$$= \alpha_R var(\bar{p}) - \alpha_R var(\bar{p})$$

$$= 0$$

[7] This approximation becomes an equality in the limit
as the length of the holding period approaches zero
since the instantaneous rate of return is by its
nature a continuously compounded rate.

[8] Proof that $\rho^2 \equiv S$:

$$\rho^2 = \frac{cov(\tilde{r}_h, p)^2}{\sigma_1^2 var(\tilde{r}_h)}$$

$$\frac{(\alpha\sigma_1^2)^2}{\sigma_1^2(\alpha^2\sigma_1^2 + \sigma_2^2)}$$

$$\frac{1}{1 + \dfrac{\sigma_2^2}{\alpha^2\sigma_1^2}} .$$

[9] For a thorough review of the capital asset pricing
model and the literature on that model see Jensen (1972).

[10] See Black (1972) for a proof of this proposition.

[11] There is a great deal of evidence that the real re-
turn on common stocks is also negatively correlated with
unanticipated inflation. See, for example, Body (1974).

[12] See Merton (1972) for a proof of this proposition.

References

K. Bhatia, "Index-Linking of Financial Contracts," Research
Report 7412, 1974, Dept. of Economics University of Western
Ontario.

F. Black, "Capital Market Equilibrium with Restricted Borrowing,"
J. of Business, 1972, 45, 444-454.

Z. Body, "Common Stocks as a Hedge Against Inflation," Boston
University School of Management Working Paper, 1974.

M. Friedman, Column, Newsweek, January 21, 1974, p. 80.

M. Jensen, "Capital Markets: Theory and Evidence," Bell Journal
of Economics and Management Science, Autumn 1974, 357-98.

M. Lovell and R. Vogel, "A CPI-Futures Market," JPE, 1971,
1009-1012.

H. Markowitz, "Portfolio Selection," Journal of Finance, 1952,
7, 77-91.

R. C. Merton, "Lifetime Portfolio Selection under Uncertainty:
The Continuous-Time Case," Rev. Econ. Statist., 1969, 51, 247-257.

_____, "Optimum Consumption and Portfolio Rules in a
Continuous-Time Model," J. of Economic Theory, 1971, 3, 373-413.

_____, "An Analytic Derivation of the Efficient Portfolio
Frontier," J. of Financial and Quant. Anal., 1972, 1851-1871.

F. Modigliani, "Some Economic Implications of the Indexing
of Financial Assets with Special Reference to Mortgages,"
Proceedings of the Conference on The New Inflation, 1974.

J. Tobin, "An Essay on Principles of Debt Management," Fiscal and
Debt Management Policies, Englewood Cliffs, New Jersey:
Prentice-Hall, 1963, 143-218.

811-75